Robust Energy-Constrained Frequency Reserves from Aggregations of Commercial Buildings

Evangelos Vrettos*[1], Frauke Oldewurtel†[2], and Göran Andersson‡[1]

[1]EEH-Power Systems Laboratory, Swiss Federal Institute of Technology (ETH) Zurich, Switzerland
[2]Department of Electrical Engineering & Computer Sciences, University of California, Berkeley, USA

June 18, 2015

Abstract

It has been shown that the heating, ventilation, and air conditioning (HVAC) systems of commercial buildings can offer ancillary services to power systems without loss of comfort. In this paper, we propose a new control framework for reliable scheduling and provision of frequency reserves by aggregations of commercial buildings. The framework incorporates energy-constrained frequency signals, which are adopted by several transmission system operators for loads and storage devices. We use a hierarchical approach with three levels: (i) reserve capacities are allocated among buildings (e.g., on a daily basis) using techniques from robust optimization, (ii) a robust model predictive controller optimizes the HVAC system consumption typically every 30 minutes, and (iii) a feedback controller adjusts the consumption to provide reserves in real time. We demonstrate how the framework can be used to estimate the reserve capacities in simulations with typical Swiss office buildings and different reserve product characteristics. Our results show that an aggregation of approximately 100 buildings suffices to meet the 5 MW minimum bid size of the Swiss reserve market.

1 Introduction

In many countries, large amounts of renewable energy sources (RES) are integrated in the power system and they increase the operational uncertainty due to their fluctuating nature. In a power system, supply and demand of electric power must be

*vrettos@eeh.ee.ethz.ch
†oldewurtel@berkeley.edu
‡andersson@eeh.ee.ethz.ch

balanced to keep frequency and voltage close to their nominal values. To maintain this balance, transmission system operators (TSOs) procure ancillary services (AS), e.g., frequency and voltage control, traditionally from conventional generators. With increasing RES shares, the need for AS is also increasing [1]. For this reason, there has been a recent interest in offering AS by flexible loads, which is known as demand response (DR). If properly aggregated, loads can provide AS without environmental impacts and, possibly, more efficiently and at a lower cost compared to generators [2].

Loads with thermal storage are suitable for DR because their consumption can be shifted in time without impact on consumer comfort. A lot of research focused on modeling, control, and estimation of large aggregations of thermostatically controlled loads (TCLs), i.e., small residential loads such as refrigerators, air conditioners, space and water heaters [3–7]. The heating, ventilation, and air conditioning (HVAC) systems of commercial buildings are also interesting for DR because they are large consumers with high thermal inertia, and they usually include a building automation and communication system that facilitates DR implementation [8]. Recently, DR applications with commercial buildings has been the topic of several papers some of which are summarized below. The role of information systems and automated model-based control for energy efficiency and DR was discussed in [9]. Commercial HVAC systems are typically complex with many control variables and cascaded loops, and thus low-order models are used for control purposes [10–12]. Research on DR with commercial buildings has focused mainly on "slow" services. For example, [13–16] proposed optimization-based approaches and model predictive control (MPC) for peak shedding or load shifting to minimize energy costs.

This paper is concerned with frequency control with commercial buildings, in particular secondary frequency control (SFC), which is also known as automatic generation control (AGC), load frequency control (LFC), automatic frequency restoration reserve (FRR), or regulation service. Recently, some authors have considered either scheduling or provision of "fast" AS from commercial buildings including SFC reserves. Ref. [17] experimentally investigated the feasibility of offering up and down regulation products with university campus buildings, and identified baseline prediction and latency as potential obstacles for implementation. The authors of [18, 19] investigated SFC via fan power control in buildings with variable air volume (VAV) systems. Based on simulations, the authors found that up to 15% of a building's fan power can be offered as reserves without significant loss of comfort, if the SFC signal is within the frequency band $f \in [1/(10 \text{ min}), 1/(4 \text{ sec})]$. In [20], a control approach similar to that of [18, 19] was experimentally validated in a real building, but without considering frequencies below $1/(10 \text{ min})$ to avoid effects on chiller power consumption. The follow-up work [21] included chiller control, which enlarged the frequency band of SFC signals to $1/(60 \text{ min})$. Ref. [22] investigated SFC by direct control of a heat pump's (HP) compressor power using a variable speed drive motor in a water-based HVAC system.

Apart from accurate tracking of the SFC signal, a TSO needs guarantees that the reserve capacity of commercial buildings will be *reliably* available. Because the buildings are energy-constrained resources reserve scheduling is required, in particular if the SFC signal is not approximately zero-energy over short periods but can be biased towards one direction over long periods of time. A priori reserve scheduling allows buildings to participate in markets for such reserve products without compromising occupants' comfort. Ref. [23] addressed this issue by developing an MPC-based method to quantify the flexibility of a commercial building and a contractual framework to declare it to the utility. In this paper, we follow this line of research and consider robust reserve scheduling for *aggregations* of commercial buildings.

The main contributions of this paper are threefold. First, we propose a new framework to estimate the SFC reserve *capacity* that can be reliably offered by an aggregation of commercial buildings considering weather conditions, occupancy, electricity prices, reserve payments, and comfort zone. The framework builds on a hierarchical control scheme with three levels, namely reserve scheduling and allocation, HVAC control, and reserve provision, and is based on robust optimization and MPC. The framework actively allocates reserves among aggregation's buildings based on their individual characteristics, which is expected to maximize the reserve potential in markets with typical requirements such as constant reserve capacity over a minimum duration and/or equal up- and down-reserve capacities. This is in contrast to [18, 19, 21] that estimated the capacity of a group of buildings by simply scaling up the estimated capacity of a single building. Second, we propose new methods to estimate reserve capacities in case of energy-constrained frequency signals, a practice adopted by several TSOs for loads or storage devices, e.g., [24]. This is a significant advantage of our method compared to [23]. A relevant approach is taken in [25], where the reliable AS capacity from a commercial building is expressed as a function of the AS signal's frequency. In contrast to [25], we are interested in the dependence of capacity on the integral of the SFC signal, i.e., on the signal's bias directly, which we believe is the major limiting factor when extracting reserves from energy-constrained resources. Third, we derive upper bounds on the untapped potential for SFC by different building types and for different reserve product characteristics such as duration, symmetry, and energy content.

Although in some cases the SFC signal is approximately zero-energy, this is not generally true for all power systems. Ideally the bias of the SFC signal would be obtained from generators, e.g., through activation of tertiary control reserves or in a nearly real-time (e.g., 5 minute) market. However, nearly real-time markets do not exist in many control areas, and activation of tertiary control reserves is not always more economical than SFC reserves, because it highly depends on the the availability of cheap and flexible generators. In addition, with increasing RES penetration in power systems, the share of conventional generators in the production mix will decrease. Despite the improvements in forecasting techniques, the RES forecasts will never be perfect and biased forecast errors will translate into biased

SFC signals. With fewer controllable generators in the production mix, other resources will have to absorb these biases. For this reason, we believe that accounting for biases in the SFC signal will be very important in the future, possibly also in power systems where the SFC signal is today approximately zero-energy. To the authors' best knowledge, this is the first time that methods are proposed to allow building aggregations to systematically schedule the reserve capacity that they can offer depending on how much bias exists in the SFC signal.

This paper includes substantial extensions over our previous work [26], where some preliminary results were presented. A major improvement is related to modeling of the SFC reserve requests within the reserve scheduling and building control optimization problems. This paper considers also energy-constrained SFC signals apart from conventional unconstrained signals, whereas [26] considered only the latter. In the absence of energy constraints, the worst case reserve request is equal to either the up- or down-reserve capacity along the whole scheduling horizon. The reserve scheduling problem for this case was cast as a robust optimization problem with additive uncertainty in [26]. On the contrary, the determination of the worst case reserve request along the scheduling horizon is more involved with energy constraints due to time coupling. In this paper, we model an energy-constrained SFC signal as an uncertain variable that lives in a polyhedral uncertainty set and enters multiplicatively in the constraints. This allows us to reformulate the robust reserve scheduling problem into a deterministic tractable one. Another main extension is related to the number and type of actuators that provide reserves in each building. In contrast to [26] that assumed either only heating or only cooling actuators per building, the proposed formulations can be used to allocate reserves among different types of actuators within the same building. Therefore, the proposed formulations are more general than the ones in [26]. In addition, the simulation studies of this paper contain new material compared with [26]. We analyze recorded data of the Swiss SFC signal to illustrate that the SFC signal is not necessarily approximately zero-energy, but it can be significantly biased over periods of several hours. The performance of the control framework is then evaluated using an energy-constrained SFC signal, whereas the sensitivity analysis is performed considering signals with and without energy constraints.

The rest of the paper is organized as follows: Section 2 summarizes some important aspects of power system SFC, Sections 3 and 4 introduce the modeling and control framework, the performance of the framework is demonstrated and evaluated in Section 5 in a simulation example, a sensitivity analysis is performed in Section 6, and Section 7 concludes the paper.

2 Frequency Control in Power Systems

Typically, a TSO controls frequency in three steps: primary, secondary and tertiary control. Primary control is a distributed, proportional controller that stabilizes the frequency after a disturbance. Secondary control is a centralized, proportional-

integral (PI) controller that restores the frequency to its nominal value and maintains the desired exchanges between neighboring control areas. Tertiary control releases secondary control in case of large disturbances and is typically manually activated.

Before continuing, we summarize some important aspects of scheduling and activation of SFC in power systems [27], which will be particularly relevant to the sensitivity analysis of Section 6. In most of Europe, SFC reserves are procured in a market setting, i.e., the generators bid their reserve capacity and price in weekly or daily auctions. The requirements of these auctions and the characteristics of the SFC reserve products vary from country to country. The minimum bid size is typically in the range $[1, 10]$ MW, e.g., 5 MW in Switzerland [28]. In many countries, only symmetric reserves, i.e., equal up- and down-reserve capacities, are allowed, whereas in other countries asymmetric reserves are also accepted. The reserve energy is requested from the generators via a signal sent by the TSO, typically every $2 - 4$ seconds. There are two main activation rules: (a) the pro-rata activation, where the reserve energy is proportional to the capacity, and (b) the merit-order activation, where the reserves are requested based on the short term marginal costs of generators. The remuneration of reserve energy is also country dependent but it is typically separate from the reserve capacity remuneration. In some countries, e.g., in Switzerland, the reserve energy remuneration is coupled with the energy price in the spot market. As a final remark, note that we treat the SFC signal as uncertain in this paper because it is unknown at the time when the reserve capacities are determined.

3 Modeling

3.1 Building Aggregation

We consider buildings with integrated room automation (IRA) systems, where heating, cooling, ventilation, blinds, and lighting are jointly controlled. IRA is typically used in office buildings because it provides high comfort while being energy efficient [29]. We represent building *thermal* dynamics using a well-tested 12^{th} order multiple-input-multiple-output bilinear model [11]. This model is based on the widespread thermal resistance-capacitance network approach, and was validated against the well-known building simulation software TRNSYS [11]. The model's satisfactory accuracy and its relatively low complexity make it suitable for MPC. A model similar to the one considered here was used in an MPC implementation in a real building, and it was found that it captures the building thermal dynamics well [10]. Furthermore, this model is flexibly customizable and allows us to perform large-scale simulation studies with different representative building types.

The original model is bilinear between inputs u and disturbances v (e.g., blind position and solar radiation), as well as between inputs and states x (e.g., blind position and room temperature). If the disturbances are fixed, e.g., according to

their predicted values, the bilinearities between u and v vanish and the system becomes time-varying. However, the bilinearities between x and u remain. For optimization purposes, sequential linear programming (SLP) can be applied by iteratively linearizing the bilinear x terms around the most recently calculated x trajectory and solving the resulting linear program (LP) until convergence [29].

If SLP is applied, the dynamics of a building b can be described by the linear time-varying model

$$x_{t+1}^b = A^b x_t^b + B_t^b u_t^b + E^b v_t^b + R^b \Delta u_t^b \,, \tag{1}$$

where $x_t^b \in \mathbb{R}^{n_x}$ denotes the states at time step t (n_x is the number of states), i.e., the room air temperature as well as the temperatures in different layers in the walls, floor, and ceiling (all measured in ^{o}C). $u_t^b \in \mathbb{R}^{n_u}$ denotes the IRA control inputs, namely heating and cooling power, ventilation, blind position, and lighting (n_u is the number of actuators). The heating and cooling are represented in the thermal model as heat fluxes affecting the system states and their units are W/m^2, i.e., the heat fluxes are normalized by the floor area. The blind position is a number between 0 (fully closed) and 1 (fully open). The lighting is also normalized by the floor area and measured in W/m^2. $v_t^b \in \mathbb{R}^{n_v}$ denotes the disturbances that affect building's states, e.g., ambient temperature in ^{o}C, solar radiation in W/m^2, and internal heat gains by occupants and equipment in W/m^2 (n_v is the number of disturbances). $\Delta u_t^b \in \mathbb{R}^{n_r}$ denotes the uncertain change in heat fluxes due to change in power consumption of the heating/cooling devices during reserve provision, where $n_r \leq n_u$ is the number of actuators that are used for reserve provision. R^b consists of the columns of B_t^b that correspond to the heating/cooling actuators that provide reserves. The system's output equation is $y_t^b = C^b x_t^b + D_t^b u_t^b + F^b v_t^b$, where $y_t^b \in \mathbb{R}^{n_y}$ denotes the room temperature in ^{o}C and illuminance in lux.

Denote by $r_{e,t}^{b,j}$ the (electric) reserve capacity of actuator j of building b at time step t. Buildings can provide *up-reserves* by decreasing their consumption, and *down-reserves* by increasing it[1]. For now, we assume symmetric reserve capacities; asymmetric reserves will be discussed in Section 4. Since the HVAC control input for heating and cooling is defined as a heat flux, it is convenient to define also the "thermal" reserve capacity $r_t^{b,j}$ that has W/m^2 units. $r_{e,t}^{b,j}$ can be obtained from $r_t^{b,j}$ by division with the coefficient of performance (COP). For notational convenience, we use the variable $r_t^{b,j}$ in the problem formulations and call it simply reserve capacity keeping in mind that it is actually the "thermal" reserve capacity. In the pro-rata activation case, the reserve energy is proportional to the reserve capacity based on a normalized SFC signal $w_t \in [-1, 1]$. Note that the reserve capacity $r_t^{b,j}$ is a decision variable for the buildings, whereas the normalized SFC signal w_t is uncertain. The primitive uncertainty w_t results in an uncertain change in electric

[1] The (electric) reserve capacity is the amount of SFC reserves that becomes available to the TSO [30]. In the context of frequency regulation, the term up-reserves denotes increase of a generator's production or decrease of a load's consumption to increase system frequency. Similarly, the term down-reserves denotes decrease of a generator's production or increase of a load's consumption to decrease system frequency.

power consumption of actuator j, $\Delta u_{e,t}^{b,j} = r_{e,t}^{b,j} w_t$ and the corresponding uncertain change in heat flux $\Delta u_t^{b,j} = r_t^{b,j} w_t$.

Denote by $x_{t+k|t}^b \in \mathbb{R}^{n_x}$ the predicted state of building b for time $t+k$ at time t. The predicted states at time t along a prediction horizon N are assembled in one vector as $\mathbf{x}_t^b = [x_{t|t}^b \;\; x_{t+k|t}^b \; ... \; x_{t+N|t}^b]^\top \in \mathbb{R}^{n_x(N+1)}$. Adopting the same notation for inputs and disturbances, the building dynamics along N can be written as $\mathbf{x}_t^b = \mathbf{A}^b x_0^b + \mathbf{B}_t^b \mathbf{u}_t^b + \mathbf{E}^b \mathbf{v}_t^b + \mathbf{R}^b \boldsymbol{\Delta u}_t^b$ and $\mathbf{y}_t^b = \mathbf{C}^b \mathbf{x}_t^b + \mathbf{D}_t^b \mathbf{u}_t^b + \mathbf{F}^b \mathbf{v}_t^b$, where the matrices $\mathbf{A}^b, \mathbf{B}_t^b$, $\mathbf{E}^b$, $\mathbf{R}^b$, $\mathbf{C}^b$, $\mathbf{D}_t^b$, and $\mathbf{F}^b$ are of appropriate dimensions. The constraints on outputs (thermal comfort zone) and HVAC control inputs along N are

$$\mathbf{y}_{\text{min}}^b \leq \mathbf{y}_t^b \leq \mathbf{y}_{\text{max}}^b \;, \;\; \mathbf{u}_{\text{min}}^b \leq \mathbf{u}_t^b + \mathbf{H}^b \boldsymbol{\Delta u}_t^b \leq \mathbf{u}_{\text{max}}^b \;, \tag{2}$$

where $\mathbf{H}^b \in \mathbb{R}^{n_u \times n_r}$ has 0 or 1 as entries. By substituting the dynamics in (2), the constraints can be written in terms of the control inputs and uncertainty as $\mathbf{G}^b \mathbf{u}_t^b + \mathbf{S}^b \boldsymbol{\Delta u}_t^b \leq \mathbf{Q}^b$, where the matrices $\mathbf{G}^b$, $\mathbf{S}^b$, and $\mathbf{Q}^b$ are defined as

$$\mathbf{G}^b = \begin{bmatrix} \mathbf{G}_p^b \\ \mathbf{I}_N \\ -\mathbf{G}_p^b \\ -\mathbf{I}_N \end{bmatrix}, \mathbf{S}^b = \begin{bmatrix} \mathbf{S}_p^b \\ \mathbf{H}^b \\ -\mathbf{S}_p^b \\ -\mathbf{H}^b \end{bmatrix}, \mathbf{Q}^b = \begin{bmatrix} \mathbf{y}_{\text{max}}^b - \mathbf{Q}_p^b \\ \mathbf{u}_{\text{max}}^b \\ \mathbf{Q}_p^b - \mathbf{y}_{\text{min}}^b \\ -\mathbf{u}_{\text{min}}^b \end{bmatrix}, \tag{3}$$

$$\mathbf{G}_p^b = \mathbf{C}^b \mathbf{B}_t^b + \mathbf{D}_t^b \;, \; \mathbf{S}_p^b = \mathbf{C}^b \mathbf{R}^b \;, \tag{4}$$

$$\mathbf{Q}_p^b = \mathbf{C}^b (\mathbf{A}^b x_0^b + \mathbf{E}^b \mathbf{v}_t^b) + \mathbf{F}^b \mathbf{v}_t^b \;, \tag{5}$$

and $\mathbf{I}_N$ is the N-dimensional identity matrix. Using $\mathbf{G}^b$, $\mathbf{S}^b$, and $\mathbf{Q}^b$, the HVAC input and comfort zone constraints along the prediction horizon can be represented compactly in the optimization problems of Section 4.

We are interested in building aggregations for two main reasons. First, the reserve markets typically have requirements on the minimum size of the bidden reserve capacity, which is typically in the range $[1, 10]$ MW, as mentioned in Section 2. In most cases, individual commercial buildings cannot meet these minimum bid size requirements, so building aggregations are needed to enable participation in the reserve market. Even in markets with low minimum bid size requirements, e.g., in the range of a few hundred kWs, building aggregations would still be of interest in presence of other typical requirements, such as symmetry and minimum duration of the bidden reserve capacity. As shown in [26], aggregating buildings with different characteristics results in a larger total reserve capacity compared with the case where each building participates individually in the market. The second argument in favor of building aggregations is more practical. An aggregator's job would be to determine the reserve capacity, bid it in the reserve market, and interact with the TSO during reserve activation and for the financial settlement. These tasks are very different to the normal activities of a building manager; therefore, the aggregator could take over this burden that would otherwise be with the building manager.

Consider an aggregation of L buildings, i.e., $b = \{1 \dots L\}$. Denote by $\mathbf{x}_t = \left[\mathbf{x}_t^1 \ \mathbf{x}_t^2 \ \dots \ \mathbf{x}_t^L\right]^\top \in \mathbb{R}^{n_x L(N+1)}$ the vector containing all predicted states of all buildings along N. Using the same notation for inputs and disturbances, the input/output constraints of the aggregation can be written as $\mathbf{G}\mathbf{u}_t + \mathbf{S}\boldsymbol{\Delta}\boldsymbol{u}_t \leq \mathbf{Q}$, where $\mathbf{G}$, $\mathbf{S}$, and $\mathbf{Q}$ are block diagonal matrices with $\mathbf{G}^b$, $\mathbf{S}^b$, and $\mathbf{Q}^b$ on the diagonal, respectively. Denoting by $\mathbf{w}_t \in [-1,1]^N$ the SFC signal along N, and by $\mathbf{r}_t \in \mathbb{R}^{n_r L N}$ a collection of the reserve capacities of all actuators and all buildings along N, the input/output constraints of the aggregation can be written as

$$\tilde{\mathbf{R}} = \begin{bmatrix} R_{1,1} & 0 & \dots & & 0 \\ R_{2,1} & R_{2,2} & 0 & \dots & 0 \\ & & \vdots & & \\ R_{N,1} & R_{N,2} & \dots & & R_{N,N} \end{bmatrix}, \mathbf{G}\mathbf{u}_t + \mathbf{S}\tilde{\mathbf{R}}\mathbf{w}_t \leq \mathbf{Q} \ , \tag{6}$$

where $\boldsymbol{\Delta}\mathbf{u}_t = \tilde{\mathbf{R}}\mathbf{w}_t$, $\tilde{\mathbf{R}} \in \mathbb{R}^{(n_r L N) \times N}$ is the *reserve capacity matrix*, and $R_{i,k} \in \mathbb{R}^{n_r L}$ is a column vector. The diagonal vectors of $\tilde{\mathbf{R}}$ are the actuators' reserve capacities for every time step t, i.e., $R_{t,t} = [r_t^1 \dots r_t^{n_r L}]^\top$. This lower triangular structure satisfies causality and allows us to model the effect of past SFC signals. For example, $R_{t,t-1}$ accounts for the signal at time $t-1$ to determine the reserve at time t. However, to comply with today's practice we fix $R_{i,k} = 0$ for $i \neq k$ in the following, i.e., the reserve energy depends only on the capacity and the current SFC signal. Thus, $\tilde{\mathbf{R}}$ is a block diagonal version of $\mathbf{r}_t$.

3.2 Uncertain Frequency Signals

3.2.1 Power Constraints (PC)

In pro-rata activation [28], the reserves are requested by the TSO via a normalized frequency signal $w_t \in [-1,1]$. This box constraint represents the power constraints of the signal. The values $w_t = -1$ and $w_t = 1$ indicate full activation of up- and down-reserves, respectively. The uncertainty set along N can be written as

$$\mathbb{W}_1 = \{\mathbf{w}_t \in \mathbb{R}^N : ||\mathbf{w}_t||_\infty \leq 1\} \ . \tag{7}$$

3.2.2 Power and Energy Constraints (PEC)

To facilitate the participation of energy-constrained resources (thermal loads and storages) in AS, the frequency signals can be distinguished by their energy content. A building aggregation could then choose the preferred product to offer reserves. Energy constraints can be cast as linear constraints on the mean value of the signal as

$$-\varepsilon T \leq \sum_{k=t}^{t+T-1} w_k \leq \varepsilon T \ , \tag{8}$$

where ε is the bias coefficient, T is the averaging period, and both are fixed by the TSO for a particular product. Equation (8) implies that the bias of the signal over T is bounded. If we stack (8) along N, we get the polyhedral constraint on the uncertainty $\mathbf{A}_w \mathbf{w}_t \leq \mathbf{b}_w$, where $\mathbf{A}_w$ is a matrix with entries -1, 0 or 1, and $\mathbf{b}_w$ is a vector with entries εT. The power constraints (7) are still present, since the full reserve capacity could be requested anytime. Denote by $\mathbf{I}_N$ the N-dimensional identity matrix, and by $\mathbf{1}_N$ the N-dimensional vector with ones. Defining $\bar{\mathbf{A}}_w = [\mathbf{A}_w; \mathbf{I}_N; -\mathbf{I}_N]$ and $\bar{\mathbf{b}}_w = [\mathbf{b}_w; \mathbf{1}_N; \mathbf{1}_N]$, the uncertainty set for PEC is

$$\mathbb{W}_2 = \{\mathbf{w}_t \in \mathbb{R}^N : \bar{\mathbf{A}}_w \mathbf{w}_t \leq \bar{\mathbf{b}}_w\} \,. \tag{9}$$

3.3 Other Sources of Uncertainty

Apart from the SFC signal, in general, there are two other sources of uncertainty associated with our problem: (a) weather and occupancy uncertainties, i.e., deviations from the predicted values, and (b) electricity and reserve price uncertainties. For simplicity, weather/occupancy uncertainties are not considered in this paper, i.e., the predictions are assumed to be perfect. There exist several approaches in the literature of building climate control to handle these uncertainties such as stochastic MPC with additive Gaussian noise [29], scenario-based MPC [31], or simply constraint tightening, which can be relatively easily integrated in our framework. We assume that the building aggregation acquires energy from the retail market, and since the utility company tariffs are typically constant over long periods, there is no electricity price uncertainty in our case. However, the situation is different for the reserve prices. Although the reserve bid price is selected by the aggregator, there is uncertainty involved because the bid might not be accepted in the auction. Incorporating this source of uncertainty requires modeling the market clearing process and is beyond the scope of this paper, but it may well constitute an interesting topic for future research.

4 Control Framework

4.1 Hierarchical Scheme

In this section, we present the hierarchical control scheme for frequency reserve provision by commercial building aggregations, which is graphically shown in Fig. 1. Level 1 (Lv1), the *aggregator scheduling*, is performed on a daily basis centrally by an aggregator. There are two conflicting objectives in the problem: (a) minimize energy consumption through energy efficient control, and (b) deviate from the energy-optimal solution to leave slack for reserve provision. The optimal tradeoff between these conflicting objectives depends on the monetary incentives for reserves, which we call capacity payments, and the cost of electricity. Lv1 determines the reserve capacity, and its allocation among buildings, which achieve

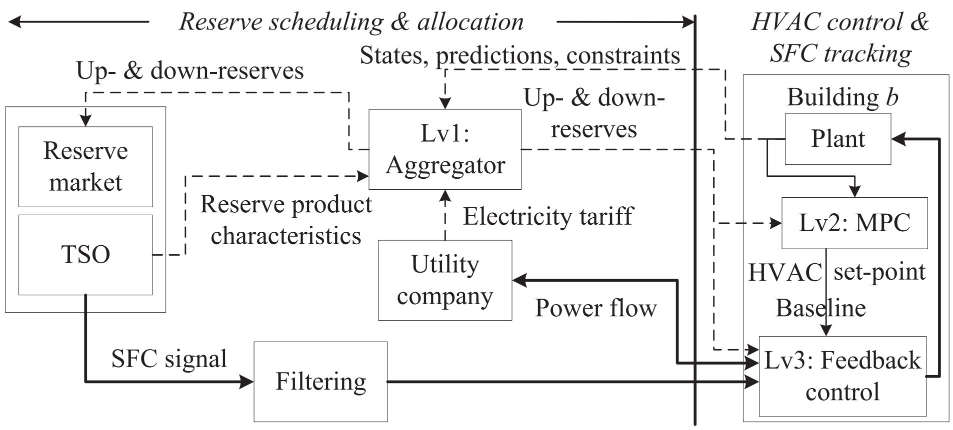

Figure 1: Overview of the proposed hierarchical control scheme.

the optimal tradeoff while respecting occupants' comfort in case the reserve is requested. The aggregator solves a robust optimization problem using building state measurements, prices of electricity and reserves, as well as predictions of weather and occupancy. We formulate the robust problem first for PC only and second for PEC. Reliable provision of frequency reserves is critical for power system security, and so TSOs typically require 100% availability of the reserve capacity [28]. For this reason, HVAC input constraints are truly hard constraints and robust optimization is a natural approach to follow.

Level 2 (Lv2), the *building HVAC control*, is a robust MPC [32] that determines the energy optimal HVAC control inputs every 30 minutes locally at each building, and leaves enough slack for reserves. Level 3 (Lv3), the *frequency signal filtering and tracking*, is a feedback controller that tracks the SFC signal by controlling the power consumption in real-time, e.g., every 10 seconds. In Fig. 1, signals and control actions with thick/solid curves are real-time, the ones with thin/solid curves are executed every 30 minutes, and the ones with dashed curves are executed once a day.

4.2 Lv1: Aggregator Scheduling

The aggregator's goal is to determine the optimal amount of reserves $\tilde{\mathbf{R}}^*$ to be offered in the market. Denote by $\mathbf{c}_t \in \mathbb{R}^{n_u L N_1}$ and $\mathbf{k}_t \in \mathbb{R}^{n_r L N_1}$ the electricity cost and reserve capacity payment vectors, respectively, where N_1 is the prediction horizon. Note that efficiency factors incorporating actuators' COP and building distribution system losses are included in $\mathbf{c}_t$ and $\mathbf{k}_t$. Lv1 can be cast as the robust

LP

$$(\mathbf{u}_t^*, \tilde{\mathbf{R}}^*) := \arg\min \ \mathbf{c}_t^\top \mathbf{u}_t - \mathbf{k}_t^\top \tilde{\mathbf{R}} \mathbf{1}_{N_1} \tag{10a}$$

$$\text{s.t. } \max_{\mathbf{w}_t \in \mathbb{W}} (\mathbf{G}\mathbf{u}_t + \mathbf{S}\tilde{\mathbf{R}}\mathbf{w}_t) \leq \mathbf{Q} \tag{10b}$$

$$\mathbf{M}\tilde{\mathbf{R}}\mathbf{1}_{N_1} = \mathbf{0} \ , \tag{10c}$$

where $\mathbb{W}$ is the uncertainty set, i.e., $\mathbb{W} \in \{\mathbb{W}_1, \mathbb{W}_2\}$. Equation (10b) requires that input and output constraints are satisfied even in the worst case of uncertainty realization. By appropriately selecting $\mathbf{M}$ in (10c), we can impose constant reserve capacities over a period of time and/or the block diagonal structure on $\tilde{\mathbf{R}}$ discussed in Section 3.1. The building dynamics in (10b) are decoupled among buildings; however, the coupling comes via (10c) and the objective function.

Denote by $\mathbf{X}_{(j)}$ the j^{th} row of any matrix or vector $\mathbf{X}$. We derive the robust counterpart of (10) for PC and PEC. Consider the j^{th} row of (10b) $\max_{\mathbf{w}_t}(\mathbf{G}_{(j)}\mathbf{u}_t + \mathbf{S}_{(j)}\tilde{\mathbf{R}}\mathbf{w}_t) \leq \mathbf{Q}_{(j)}$. The term $\mathbf{S}_{(j)}\tilde{\mathbf{R}}\mathbf{w}_t$ is a scalar and can be written as $\mathbf{w}_t^\top \mathbf{F}_{(j)}\tilde{\mathbf{r}}$, where $\tilde{\mathbf{r}} \in \mathbb{R}^{n_r L N_1^2}$ is a column-wise vectorized version of $\tilde{\mathbf{R}}$ and $\mathbf{F}_{(j)} = (\mathbf{I}_{N_1} \otimes \mathbf{S}_{(j)}^\top)^\top \in \mathbb{R}^{N_1 \times (n_r L N_1^2)}$, where $\otimes$ is the Kronecker product. Thus, the j^{th} row of (10b) is equivalent to

$$\max_{\mathbf{w}_t \in \mathbb{W}} (\mathbf{G}_{(j)}\mathbf{u}_t + \mathbf{w}_t^\top \mathbf{F}_{(j)}\tilde{\mathbf{r}}) \leq \mathbf{Q}_{(j)} \ . \tag{11}$$

4.2.1 Robust Counterpart for Power Constraints

In the presence of PC only, the uncertainty set $\mathbb{W}_1$ is given by (7), i.e., $\mathbf{w}_t$ is constrained in an ∞-norm ball. In this case, we can maximize the left-hand side of (11) analytically using the dual of ∞-norm, i.e., the 1-norm [33]. Following this procedure, the deterministic equivalent of (11) is $\mathbf{G}_{(j)}\mathbf{u}_t + ||\mathbf{F}_{(j)}\tilde{\mathbf{r}}||_1 \leq \mathbf{Q}_{(j)}$. Repeating this procedure for all rows of (10b), the robust counterpart problem of (10) can be written as

$$(\mathbf{u}_t^*, \tilde{\mathbf{r}}^*) := \arg\min \ \mathbf{c}_t^\top \mathbf{u}_t - \tilde{\mathbf{k}}_t^\top \tilde{\mathbf{r}} \tag{12a}$$

$$\text{s.t. } \mathbf{G}_{(j)}\mathbf{u}_t + ||\mathbf{F}_{(j)}\tilde{\mathbf{r}}||_1 \leq \mathbf{Q}_{(j)} \ \ \forall j \tag{12b}$$

$$\tilde{\mathbf{M}}\tilde{\mathbf{r}} = \mathbf{0} \ , \tag{12c}$$

where $\tilde{\mathbf{k}}_t \in \mathbb{R}^{n_r L N_1^2}$ and $\tilde{\mathbf{M}}$ are defined by $\tilde{\mathbf{k}}_t^\top \tilde{\mathbf{r}} = \mathbf{k}_t^\top \tilde{\mathbf{R}}\mathbf{1}_{N_1}$ and $\tilde{\mathbf{M}}\tilde{\mathbf{r}} = \mathbf{M}\tilde{\mathbf{R}}\mathbf{1}_{N_1}$. Problem (12) is an LP, but it grows quadratically in N_1 since we need N_1 auxiliary variables and $2N_1$ additional constraints for each uncertain constraint (11) to model the 1-norm. A similar reformulation for a general class of linear systems with reserve demands can be found in [34].

We now consider the case where either only heating or only cooling actuators provide reserves in each building, which is likely in practice to avoid energy dumping by simultaneous heating and cooling. Recall from (3), (4) that $\mathbf{S}$ is a

block diagonal collection of $\mathbf{S}^b = [\mathbf{C}^b\mathbf{R}^b; \mathbf{H}^b; -\mathbf{C}^b\mathbf{R}^b; -\mathbf{H}^b]$. If only heating (cooling) actuators are used for reserve provision, then all entries of $\mathbf{R}^b$ are non-negative (non-positive). Additionally, all entries of $\mathbf{C}^b$ and $\mathbf{H}^b$ are non-negative by construction. Therefore, every row of $\mathbf{S}$ contains either only non-negative or only non-positive entries and, by construction, the same holds for all entries of $\mathbf{F}_{(j)}$, i.e., $[\mathbf{F}_{(j)}(i,k) \geq 0 \; \forall i,k]$ or $[\mathbf{F}_{(j)}(i,k) \leq 0 \; \forall i,k] \; \forall j$. Based on the definitions of $\mathbf{F}_{(j)}$, $\mathbf{S}_{(j)}$, $\tilde{\mathbf{r}}$, and $\mathbf{r}_t$, and recalling that $\tilde{\mathbf{r}}$ is non-negative, $||\mathbf{F}_{(j)}\tilde{\mathbf{r}}||_1$ can be equivalently written as the linear term

$$||\mathbf{F}_{(j)}\tilde{\mathbf{r}}||_1 = \sum_{i=1}^{N_1} \sum_{k=1}^{n_{\text{el}}} |\mathbf{F}_{(j)}(i,k)\tilde{\mathbf{r}}(k)| = \tag{13}$$
$$\sum_{i=1}^{N_1} \sum_{k=1}^{n_{\text{el}}} |\mathbf{F}_{(j)}(i,k)|\tilde{\mathbf{r}}(k) = \mathbf{1}_{N_1}|\mathbf{F}_{(j)}|_\star \tilde{\mathbf{r}} = |\mathbf{S}_{(j)}|_\star \mathbf{r}_t,$$

where $n_{\text{el}} = n_r L N_1^2$ and $|\cdot|_\star$ denotes the element-wise absolute value operator. In this case, the more general formulation (12) can be simplified to the following LP that has the same size as (10), and so can be solved efficiently[2].

$$(\mathbf{u}_t^*, \mathbf{r}_t^*) := \arg\min \; \mathbf{c}_t^\top \mathbf{u}_t - \mathbf{k}_t^\top \mathbf{r}_t \tag{14a}$$
$$\text{s.t. } \mathbf{G}_{(j)}\mathbf{u}_t + |\mathbf{S}_{(j)}|_\star \mathbf{r}_t \leq \mathbf{Q}_{(j)} \;\; \forall j \tag{14b}$$
$$\mathbf{M}\mathbf{r}_t = \mathbf{0} \, . \tag{14c}$$

4.2.2 Robust Counterpart for Power and Energy Constraints

With PEC, the uncertainty set $\mathbb{W}_2$ is a polyhedron and duality theory can be applied to derive the robust counterpart problem [33, 35]. We write constraint (11) as an optimization problem over $\mathbf{w}_t$

$$\max_{\mathbf{w}_t} \quad (\mathbf{G}_{(j)}\mathbf{u}_t + \mathbf{w}_t^\top \mathbf{F}_{(j)}\tilde{\mathbf{r}}) \leq \mathbf{Q}_{(j)} \tag{15a}$$
$$\text{s.t.} \quad \bar{\mathbf{A}}_w \mathbf{w}_t \leq \bar{\mathbf{b}}_w \, . \tag{15b}$$

By deriving the dual problem of (15) for all j, we obtain the following robust counterpart of (10)

$$(\mathbf{u}_t^*, \tilde{\mathbf{r}}^*, \boldsymbol{\lambda}_{(j)}^*) := \arg\min \; \mathbf{c}_t^\top \mathbf{u}_t - \tilde{\mathbf{k}}_t^\top \tilde{\mathbf{r}} \tag{16a}$$
$$\text{s.t. } \bar{\mathbf{b}}_w^\top \boldsymbol{\lambda}_{(j)} + \mathbf{G}_{(j)}\mathbf{u}_t - \mathbf{Q}_{(j)} \leq 0 \;\; \forall j \tag{16b}$$
$$\bar{\mathbf{A}}_w^\top \boldsymbol{\lambda}_{(j)} = \mathbf{F}_{(j)}\tilde{\mathbf{r}} \;\; \forall j \tag{16c}$$
$$\boldsymbol{\lambda}_{(j)} \geq \mathbf{0} \; \forall j, \tag{16d}$$
$$\tilde{\mathbf{M}}\tilde{\mathbf{r}} = \mathbf{0} \, , \tag{16e}$$

[2] Note that (14) is a special case of (12), and it is equivalent to formulation (13)-(15) in our previous work [26].

where $\boldsymbol{\lambda}_{(j)}$ is the vector of dual variables. Problem (16) is an LP, but its size grows quadratically in N_1. Although this dualization technique holds also for PC, we apply analytic maximization in that case since it results in fewer variables and constraints.

Remark: Asymmetric reserves can be modeled defining $\tilde{\mathbf{R}}^{+}$, $\mathbf{w}_t^{+} \in [0,1]^{N_1}$ for down- and $\tilde{\mathbf{R}}^{-}$, $\mathbf{w}_t^{-} \in [-1,0]^{N_1}$ for up-reserves. In case of PC only, the uncertainty set remains a polyhedron and a tractable robust counterpart problem can be derived using analytic maximization. The reader is referred to [26] for detailed formulations. In case of PEC, the non-linear constraint $w_t^{-} w_t^{+} = 0$ is needed to ensure that up- and down- reserves are not requested simultaneously. Therefore, the uncertainty set is not a polyhedron any more and the dualization technique cannot be applied in this case to derive the robust counterpart problem.

4.3 Lv2: Building HVAC Control

Given the optimal reserve allocation from Lv1, in Lv2 the HVAC control inputs are determined every 30 minutes by the robust MPC with prediction horizon N_2

$$\mathbf{u}_t^{b,*} := \arg\min \ \left(\mathbf{c}_t^b\right)^{\top} \mathbf{u}_t^b \tag{17a}$$

$$\text{s.t.} \max_{\mathbf{w}_t \in \mathbb{W}} \left(\mathbf{G}^b \mathbf{u}_t^b + \mathbf{S}^b \tilde{\mathbf{R}}^{b,*} \mathbf{w}_t\right) \leq \mathbf{Q}^b \ , \tag{17b}$$

where $\mathbf{c}_t^b$ and $\tilde{\mathbf{R}}^{b,*}$ are the parts of $\mathbf{c}_t$ and $\tilde{\mathbf{R}}^{*}$, respectively, that are relevant for building b. The first input of the optimal control sequence of (17) determines the Lv2 setpoint of the HVAC system for the next 30 minutes[3], $u_t^{b,\mathrm{Lv2}}$. Problem (17) formulates an MPC with open-loop predictions, i.e., the optimization is performed explicitly over the control inputs $\mathbf{u}_t^b$. MPC with closed-loop predictions, i.e., optimization over affine policies of the uncertainty, showed minor or zero performance improvement in this case, and so it is not used. In the following, we derive the robust counterparts of (17) for PC and PEC.

4.3.1 Robust Counterpart for Power Constraints

In this case, the deterministic equivalent of (17) can be obtained by substituting (17b) with $\mathbf{G}_{(j)}^b \mathbf{u}_t^b \leq \mathbf{Q}_{(j)}^b - ||\mathbf{F}_{(j)}^b \tilde{\mathbf{r}}^{b,*}||_1 \ \forall j$, where $\mathbf{F}_{(j)}^b$ is defined similarly to $\mathbf{F}_{(j)}$ but for a single building, and $\tilde{\mathbf{r}}^{b,*}$ is a column-wise vectorized version of $\tilde{\mathbf{R}}^{b,*}$. Note that $\tilde{\mathbf{r}}^{b,*}$ is fixed in Lv2, and the right-hand side of the inequality is a constant.

[3] Although a consumption schedule is calculated in Lv1, the MPC of Lv2 can reduce the costs due to less uncertainty (recent reserve requests are known and better weather forecasts might be available) and, possibly, shorter optimization time steps. In case of plant-model mismatches, MPC additionally reduces constraint violations due to its closed-loop operation. The MPC schedule is the building's baseline consumption, and is communicated to the aggregator. Since the baseline is a by-product of the formulation, baseline prediction methods that have proven to be hard [17] are not required. Furthermore, the predictive nature of MPC inherently accounts for rebound effects due to reserve provision when calculating future HVAC setpoints [9].

4.3.2 Robust Counterpart for Power and Energy Constraints

With PEC, the robust counterpart of (17) is

$$(\mathbf{u}_t^{b,*}, \boldsymbol{\lambda}_{(j)}^{b,*}) := \arg\min\, (\mathbf{c}_t^b)^\top \mathbf{u}_t^b \tag{18a}$$

$$\text{s.t. } \bar{\mathbf{b}}_{w,t}^\top \boldsymbol{\lambda}_{(j)}^b + \mathbf{G}_{(j)}^b \mathbf{u}_t^b - \mathbf{Q}_{(j)}^b \leq 0 \quad \forall j \tag{18b}$$

$$\bar{\mathbf{A}}_{w,t}^\top \boldsymbol{\lambda}_{(j)}^b = \mathbf{F}_{(j)}^b \tilde{\mathbf{r}}^{b,*}\, \forall j, \tag{18c}$$

$$\boldsymbol{\lambda}_{(j)}^b \geq \mathbf{0}\, \forall j\ . \tag{18d}$$

Problem (18) is similar to (16), but there are two main differences: first, $\tilde{\mathbf{r}}^{b,*}$ is fixed; and second, $\bar{\mathbf{A}}_{w,t}$ and $\bar{\mathbf{b}}_{w,t}$ are time-varying. For a time step t in the averaging interval $[t_1, t_2]$ of length T, (8) can be written as $-\varepsilon T - w_{\mathrm{p},t} \leq \sum_{k=t}^{t_2} w_k \leq \varepsilon T - w_{\mathrm{p},t}$, where $w_{\mathrm{p},t} = \sum_{k=t_1}^{t-1} w_k$ is known because the uncertainty up to $t-1$ is realized. Thus, the coupling constraint on $\{w_t, \ldots, w_{t_2}\}$ depends on the energy content of the SFC signal in the previous time steps of the averaging interval.

4.4 Lv3: Frequency Signal Filtering and Tracking

In Lv3, the HVAC consumption is controlled around $u_t^{b,\mathrm{Lv2}}$ for reserve provision. We consider water-based HVAC systems that are common in Europe, but the proposed reserve scheduling problem applies also to air-based systems [23]. The power consumption of water circulation pumps in water-based HVAC systems is typically small, and so we directly control the heating or cooling devices, e.g., heat pumps (HPs), to provide reserves. The desired power consumption of HP j of building b is

$$u_t^{b,j,\mathrm{Lv3}} = u_t^{b,j,\mathrm{Lv2}} + \Delta u_t^{b,j} = u_t^{b,j,\mathrm{Lv2}} + w_t r_t^{b,j,*}\ , \tag{19}$$

where $r_t^{b,j,*}$ is fixed from Lv1. In case of asymmetric reserves, $r_t^{b,j,*}$ is equal to the down-reserve capacity if $w_t \geq 0$, and equal to the up-reserve capacity if $w_t < 0$. In a fast time scale, and depending on the HVAC system, the HP consumption can be controlled by modifying either the water temperature setpoint at condenser's outlet or the refrigerant's flow rate via valves. We rely on the second approach that was experimentally shown to be able to track fast reference signals, e.g., SFC signals, in [36]. The desired HP consumption $u_t^{b,j,\mathrm{Lv3}}$ can be tracked using a feedback PI controller.

With PC only, the normalized reserve request w_t sent from the TSO is the original SFC signal, whereas with PEC, w_t is a filtered version of that signal. With reasonable modifications, the current operational paradigm at the TSO side can integrate multiple SFC signals with different energy contents. In the current paradigm, there is a single SFC signal that is the output of a PI controller with the area control error (ACE) as input. In the new paradigm, the TSO provides a number of reserve products and each reserve provider chooses the product to offer

its reserve capacity. Assume that the total reserve request (original SFC signal) at time step t is $\tilde{w}_t$. The TSO will decompose $\tilde{w}_t$ into the desired number of signals (e.g., using a filter bank) and will sent the appropriate components to the providers depending on the reserve product they offer reserves for. To simulate this process in this paper, we use a causal Chebyshev filter to get the energy-constrained signal w_t sent to the building aggregation, but other filters might also be used. Similar filtering approaches have been used in previous works on power system applications, e.g., [37]. Note that although w_t is a filtered signal, it does not mean that its bias is zero but instead that its bias is bounded. The filter's transfer function is

$$H(z) = \frac{\sum_{i=0}^{n_\mathrm{f}} b_i z^{-i}}{1 + \sum_{i=1}^{n_\mathrm{f}} a_i z^{-i}} \ , \tag{20}$$

where n_f is the filter's order, and a_i, b_i are its coefficients that depend on the pass-band edge frequency f_c. From a TSO perspective, n_f and f_c can be chosen such that the resulting low-frequency component (LF signal) and high-frequency component (HF signal) have the desired ramping rates and energy contents. In this paper, we fix $n_\mathrm{f} = 3$ since it showed good performance in preliminary simulations.

It is important to clarify how the SFC signal filtering is taken into account in the higher levels of the control hierarchy. Recall that the filtered signal is tracked every few seconds in Lv3, whereas the decisions in Lv1 (reserve scheduling) and Lv2 (determination of optimal building setpoints) are taken every 30 minutes. Due to this time scale separation, the important information from the filtering of Lv3 that needs to be conveyed to Lv1 and Lv2 is only the integral of the SFC signal over this period, i.e., the bias of the signal. This is needed for example in Lv1 to schedule the reserve capacities without violating the buildings' thermal comfort constraints. By formulating the energy constraint of the SFC signal as the linear inequalities (8), we can directly account for the signal's bias in the optimization problems of Lv1 and Lv2 in a tractable way. Note that the bias coefficients ε for different averaging periods T in (8) can be empirically obtained from the filter (20). To do so, one can simply apply the filter on historical data of SFC signals using different $f_\mathrm{c} = 1/T$.

Due to the robust design, any admissible reserve request w_t that satisfies (8) will not lead to comfort zone or input constraint violations, provided that there is no significant plant-model mismatch. The tracking quality of w_t depends on HP's mechanical delays and dead-times, ramping limits, and minimum down-times and/or run-times. If such dynamics are significant, tests similar to the ones in [36] must be performed to identify upper limits on the frequency content of w_t that result in good tracking by the HP. In this case, the building aggregation could form a coalition with faster resources, e.g., an aggregation of residential electric water heaters, exclude very high frequency components of the SFC signal using an appropriate band-pass filter, and send them to the faster resources. Since the focus of this paper is on the robust reserve scheduling and allocation side, we assume that such HP dynamics are negligible, i.e., the reference $u_t^{b,j,\mathrm{Lv3}}$ (and so the SFC

Table 1: Bias Coefficients (ε) for the SFC Signal and its High-frequency Component (HF) for Different Averaging Periods (T)

T	1 h	2 h	4 h	6 h	8 h	12 h
SFC 2009	1.000	0.989	0.952	0.796	0.592	0.505
HF 2009	0.528	0.467	0.337	0.273	0.384	0.229
SFC 2012	0.927	0.781	0.674	0.624	0.553	0.448
HF 2012	0.382	0.300	0.317	0.290	0.237	0.203

signal) can be perfectly tracked in our simulations if the comfort zone and input constraints are satisfied.

5 Performance of Control Framework

5.1 Investigation Setup

We investigate the performance of the proposed control framework in simulations with an aggregation of typical Swiss office buildings. We consider two HVAC systems: in system A, heating is performed via radiators with coefficient of performance (COP) equal to 3, whereas cooling with cooled ceilings (COP $= 3.5$); in system B, both heating and cooling are performed using thermally activated building systems (TABS) with COP $= 3.4$. We also differentiate between heavy (eh) or light (el) building envelope, high (wh) or low (wl) window area fraction, and high (gh) or low (gl) internal gains. In our simulations, we consider an aggregation of 6 large buildings (15000 m^2 each) with the following characteristics: A1 = $\{$A,eh,wh,gh$\}$, A2 = $\{$A,eh,wl,gl$\}$, A3 = $\{$A,el,wl,gl$\}$, B1 = $\{$B,eh,wh,gh$\}$, B2 = $\{$B,eh,wl,gl$\}$, and B3 = $\{$B,el,wl,gl$\}$. Typical occupancy profiles were used, whereas weather data were provided by Meteoswiss (the Swiss federal office of meteorology and climatology). More information regarding the buildings can be found in [26, 29]. Because of the thermal inertia, heating/cooling actuators can be used to provide frequency reserves. In buildings with the HVAC systems considered here, the fresh air flow rate is usually kept constant because changes would be immediately realized by the occupants. For this reason, we do not use ventilation for reserve provision. The temperature comfort zone during working hours is $[21, 24]^o$C in winter and $[22, 25]^o$C in summer, which is in accordance with the ASHRAE 55-2013 standard [38]. During non-working hours and weekends, the comfort zone is relaxed to $[12, 35]^o$C in both seasons. The optimizations are performed with a time step of 30 minutes, which is the discretization step of building models (1), and the prediction horizons of Lv1 and Lv2 are fixed to $N_1 = 96$ (2 days) and $N_2 = 48$ (1 day), respectively. We assume symmetric, daily reserves, i.e., constant reserve capacity over a day, and capacity payments 10% higher than the electricity price, i.e., $k = 1.1c$.

5.2 Parameters of Energy Constraints

To apply energy constraints as in (8), we determine appropriate pairs of averaging period and bias coefficient (T, ε) based on the historical normalized SFC signals $w_k \in [-1, 1]$ from the Swiss control area for 2009 and 2012. We consider six different averaging periods, $T = \{1, 2, 4, 6, 8, 12\}$ hours, and calculate six sets of filter parameters in (20), one for each of the pass-band edge frequencies $f_\text{c} = 1/T$. For each value of T, we filter the historical SFC signals using the corresponding filter (20) to get the HF signals for 2009 and 2012. For each of the four signals and for each value of T, we calculate ε as the largest absolute mean value of the SFC or HF signal over any period T, i.e., $\varepsilon = \max(|(1/T) \cdot \sum_{k=1}^{T} w_k|)$. The results are summarized in Table 1. Notice that the original SFC signals can be significantly biased over periods of several hours. Note also that the biases of the HF signals are significantly lower than those of the original SFC signals, and that the signals in 2012 are generally less biased than in 2009. This is because from March 2012 the ACE of Switzerland is netted with that of other European countries before generating the SFC signal. We use the HF signal of 2012 and select $(T = 2\text{ h}, \varepsilon = 0.3)$ for the simulations of this section, according to Table 1.

5.3 Operation for Typical Weeks

We present simulation results for typical weeks in winter and summer. In Fig. 2, we show a 4-hour extract of the SFC, LF, and HF signals for a sample day, whereas in Fig. 3 we show the 2-hour moving averages of the signals for the whole day. Although the original SFC signal is mostly negative, its bias is absorbed by the LF signal, and so the HF signal is approximately zero-mean. The bias ε of the SFC signal is larger than 0.3, whereas the energy-constrained HF signal has a bias less than 0.15.

Figure 4 shows the optimal total reserve capacity of the aggregation and its allocation among buildings for the winter week. The capacity is constant for each day and ranges from approximately 260 kW on Friday to 380 kW on Saturday with an average weekly value of 313 kW. Note that the capacity is shifted among buildings in a way that maximizes the total capacity of the aggregation. Interestingly, the buildings offer higher reserve capacities when they normally consume less power. For example, buildings with system A contribute mainly at night, whereas buildings with system B offer more reserves during working hours, because they prefer to preheat at night. The scheduled total energy consumption of the aggregation is spread throughout the whole week to maximize the reserve potential. During the weekend, the buildings are not occupied and the comfort zone is larger, which results in higher reserve capacities compared to working days.

The optimal operation of building A3 based on Lv2 and Lv3 is presented in Figs. 5 and 6. Figure 5 shows the HVAC system heating power and Fig. 6 the resulting temperature trajectories. The black-solid curves indicate the scheduled consumption and temperature by the MPC, whereas the grey-dashed curves corre-

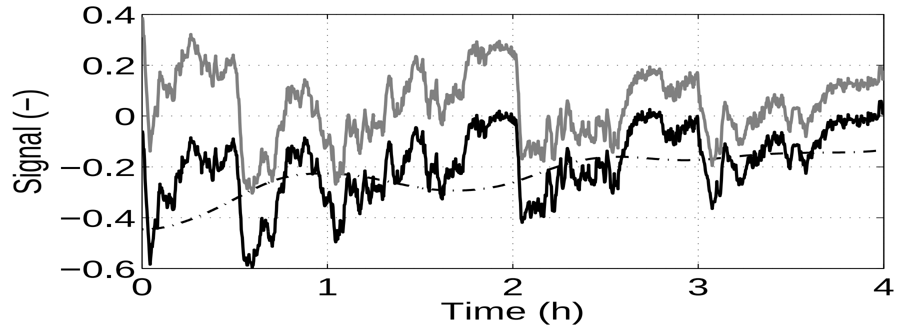

Figure 2: A 4-hour extract of the original SFC signal, its low-frequency (LF), and its high-frequency (HF) components for a sample day.

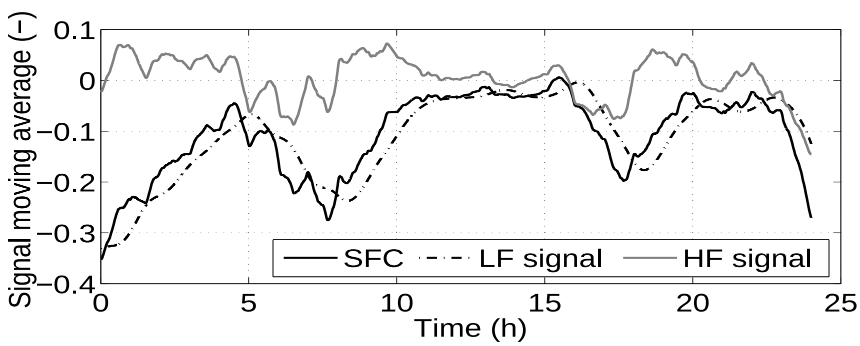

Figure 3: The moving averages of the signals of Fig. 2 for the whole sample day.

spond to the final values after tracking the HF signal. The grey envelope shows the problem's robust region, which is defined as the region of the thermal comfort zone with the following property. If the temperature is within the robust region at time step $t = 0$, then for any SFC signal with the appropriate energy content there exists a feasible control input trajectory that satisfies the thermal comfort and HVAC input constraints (the black-dashed curves in both figures) along the prediction horizon $t \in [0, N-1]$. In our simulations, the HVAC inputs and room temperature stay always within the robust region.

In summer, the reserve potential from cooling actuators is also significant. During the considered week, the capacities range from approximately 60 kW on Sunday to 420 kW on Saturday with an average weekly value of 323 kW. We observed a more uniform distribution of reserves among different buildings in summer, compared to winter, because cooled ceilings and TABS have similar time constants. Due to space limitations, plots similar to Figs. 4, 5 and 6 are omitted.

Based on our results, the considered 6 buildings with floor area 15000 m^2, average rated heating power 27 W/m^2, and average rated cooling power 32 W/m^2

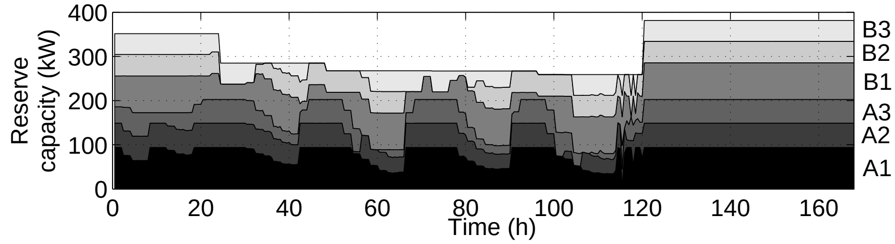

Figure 4: Optimal reserve allocation among buildings (winter week).

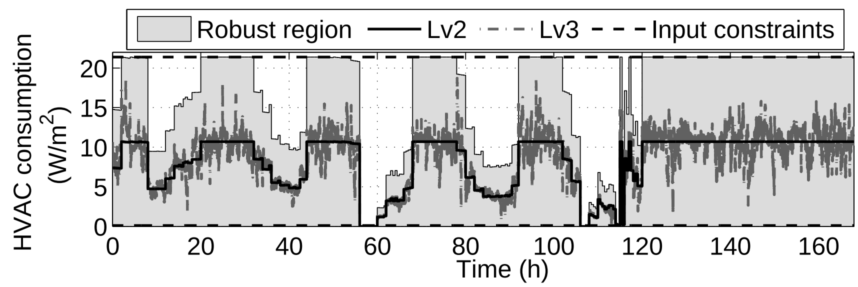

Figure 5: HVAC heating power for building A3 (winter week).

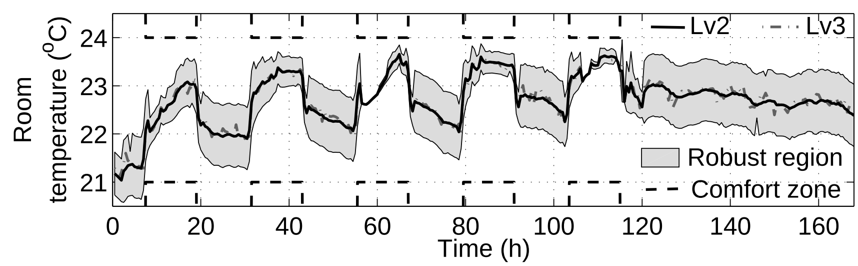

Figure 6: Room temperature for building A3 (winter week).

provide a reserve capacity equal to 313 kW and 323 kW on average in winter and summer, respectively. Using simple linear extrapolation, we find that the required minimum reserve capacity of 5 MW in Switzerland can be provided by 96 (similar) buildings in winter and 93 buildings in summer, i.e., approximately 100 buildings in both seasons. Note that this is only a rough estimate as: (a) it is based on the average reserve capacity values; and (b) the extrapolation might overestimate the required number of buildings to meet the 5 MW limit, because the larger the aggregation the more flexibility exists in allocating reserves among buildings.

5.4 Effect on Energy Consumption

To maximize the reserve potential, the buildings operate in a less energy efficient way. For example, during the considered winter week reserve provision resulted in a consumption 60% higher than that of energy efficient building control, i.e.,

an MPC with objective to minimize electricity cost without offering reserves. As shown in Fig. 5, in order to provide reserves in both directions, the buildings try to operate close to the middle of the heating/cooling device's power range. This is in contrast to an energy efficient operation, where the power consumption would be as close as possible to the minimum value. However, the increase in energy consumption does not mean that the building control is suboptimal. For the given electricity price and capacity payment, this building operation minimizes the total cost defined as the sum of electricity cost and reserve profit.

It is important to note that the 60% consumption increase is in comparison with an MPC-based energy efficient building control. However, most buildings today operate with supervisory rule-based controllers, i.e., not optimal controllers. Therefore, if the proposed methods are applied to such buildings, the observed increase in consumption is expected to be less than the reported value. Although energy efficiency is usually the goal in building control, increasing the energy consumption is not necessarily a drawback if this helps accommodate more RES in the power system, and at the same time the additional electric energy is stored as thermal energy in the buildings.

5.5 Discussion on Imperfect Disturbance Predictions

As mentioned in Section 3.3, the formulations and simulations in this paper assumed perfect weather and occupancy predictions. The robust formulation guarantees satisfaction of the HVAC and comfort constraints for any admissible reserve request and perfect weather and occupancy predictions, but it cannot provide mathematical guarantees on the worst case temperature deviations in case of imperfect predictions. Empirical guarantees could be provided by Monte Carlo simulations: (i) schedule the reserve capacities using imperfect disturbance predictions; (ii) simulate the building operation under disturbance uncertainty, i.e., the weather and occupancy realizations are different to the predictions; and (iii) analyze the results to keep track of the number and magnitude of temperature deviations. Alternatively, probabilistic guarantees could be obtained by modeling the weather and occupancy uncertainty via scenarios and then robustifying against the reserve uncertainty separately for each scenario. This is an exciting research direction for future work.

However, some intuition can be provided without following any of the above two approaches. Consider a building in energy efficient operation using a deterministic MPC that relies on an imperfect weather forecast. As explained in Section 5.4, the building would operate close to the minimum power consumption, and thus the temperature trajectory would stay close to one of the comfort zone boundaries. Now consider the same building operated under reserve provision. As shown in Fig. 6, the building would operate closer to the middle of the comfort zone to maximize the reserve capacity. For this reason, we expect the reserve provision case to result in smaller and perhaps less frequent comfort zone violations compared to the energy efficient control case, if the building is exposed to the same weather and occupancy uncertainty in both cases.

6 Sensitivity Analysis

6.1 Capacity Payments

As explained in Section 4.1, the proposed reserve scheduling methods identify the optimal tradeoff between minimizing energy consumption and leaving enough slack for reserve provision. The buildings would not deviate from the energy efficient control and would not offer any reserves if the additional electricity cost occurring due to this deviation were higher than the reward received for the slack provided as reserve capacity. In principle, the amount of reserves depends on the relationship (ratio) between the capacity payment k, i.e., the remuneration for each kW of reserve capacity provided, and the electricity price c, i.e., the cost for each kWh of electricity consumed. In this Section, we investigate this relationship by running simulations, similar to the ones in Section 5.3, over 2-week periods in winter and summer and for various k/c ratios. The total reserve capacity for each case, i.e., the sum of the capacities of each day of the 2-week period, is presented in Fig. 7, where the left plot is for ratios $k/c > 1$ and the right plot is for ratios $k/c \leq 1$. These plots represent the aggregation's bid curves because they communicate how much capacity the aggregation is willing to bid in the reserve market depending on the payment it receives for each kW of the capacity.

For $k/c > 1$ (left plot - Fig. 7a) the simulations are performed for winter and summer with and without consideration of energy-constrained SFC signals: the black curves correspond to winter weeks ("win") and the grey curves to summer weeks ("sum"), whereas the dashed curves are for PC and the solid curves for PEC. Our simulations show that with the same financial incentive and for both seasons, the buildings are willing to offer up to 10% more reserves compared with PC if energy constraints are considered. Note that the gap between PC and PEC is generally larger for lower k/c, particularly in winter. In winter, the capacity saturates at its maximum value at $k/c = 1.1$, whereas in summer it increases monotonically as the ratio increases up to 2.

The analysis of Fig. 7a focused on $k/c > 1$, which is a necessary condition for reserve provision with PC. This observation was also made in our previous work [26]; however, here we provide an explanation by studying the structure of problem (12). If $k < c$, the optimal solution is $(\mathbf{u}_t^*, \tilde{\mathbf{r}}^*) = (\mathbf{u}_t^{\text{min}}, \mathbf{0})$, where $\mathbf{u}_t^{\text{min}}$ is the energy optimal scheduling. If $k = c$, the optimal cost is 0 and any solution within the feasible range of $\tilde{\mathbf{r}}$ will be optimal. If $k > c$, the optimal solution is $(\tilde{\mathbf{r}}^{\text{max}}, \tilde{\mathbf{r}}^{\text{max}})$, where $\tilde{\mathbf{r}}^{\text{max}}$ is the upper limit of $\tilde{\mathbf{r}}$, and the optimal cost will be $(\mathbf{c}_t^\top - \mathbf{k}_t^\top)\tilde{\mathbf{r}}^{\text{max}} < 0$, i.e., the aggregation earns profit. The limit $\tilde{\mathbf{r}}^{\text{max}}$ depends on input/output constraints, and so different solutions are obtained for different $k/c > 1$ ratios, as shown in Fig. 7a. In case of daily reserves, $k/c > 1$ need not to be satisfied point-wise throughout the whole day; instead, $k/c < 1$ can be chosen during daytime, and $k/c > 1$ at night when electricity prices might be lower. Reserve provision will be triggered if $||k||_1/||c||_1 > 1$, i.e., the capacity payment is on average higher than the electricity price.

On the other hand, with energy-constrained SFC signals (PEC) reserves can be provided also with ratios $k/c < 1$ throughout the whole day. It is easier to explain this with an example. Assume that the buildings have declared a capacity $r(t)$ for day d. Assume also that up regulation (i.e., consumption decrease) is mainly requested during day d. If the signal is energy constrained, only a fraction of the worst case reserve energy $\int_{t=0}^{N-1} r(t)dt$ will be requested as consumption decrease from the buildings. The rest part of $\int_{t=0}^{N-1} r(t)dt$ will be stored as thermal energy in the buildings and will reduce the required heating/cooling energy (and the respective costs) during day $d+1$. For this reason, the buildings are willing to provide reserves even if $k/c < 1$. We present simulation results for PEC in winter and summer in Fig. 7b. The threshold ratio for reserve provision is $k/c = 0.32$ for both winter and summer, and it depends on ε and T. The capacity increases slowly in the ratio range $[0.32 - 0.99]$, particularly in summer, and then it suddenly jumps to higher values as k/c approaches to 1. No results are shown for PC in Fig. 7b because the capacity is zero for $k/c < 1$, as explained before.

Fig. 7a shows that for unconstrained SFC signals (PC) a ratio $k/c = 1.01$, which means a capacity payment 1% higher than the electricity price, taps most of the reserve potential. In particular, further increasing the ratio up to the maximum considered value $k/c = 2$ increases the reserves only by 8% in winter and no more than 18% in summer. Assuming an average electricity price of 146.6 CHF/MWh, which is the case for consumers who consume more than 60 MWh/year in Zurich, with a ratio $k/c = 1.01$ capacity payments around 148 CHF/MW/h are needed. This is significantly lower than the most expensive accepted bids, but approximately 4 times higher than the average capacity payment in 2013 [39]. Fig. 7b shows that energy-constrained SFC signals can reduce the necessary capacity payments down to 32% of the retail price ($k/c = 0.32$), but of course with a large reduction in the reserve capacity. Thus, reserves are actually costly for buildings already equipped with MPC for *energy efficient* (optimal) control, especially if the SFC signal is not energy-constrained. This is in contrast to TCLs with simple hysteresis control based on a deadband, where reserves can be provided at a lower cost. However, note that our calculations are based on the prevailing case where the buildings acquire energy in the retail electricity market. In another market setting where the buildings acquired energy directly in the spot market, the buildings could offer reserves at more competitive prices because the retail electricity prices are typically significantly higher than the wholesale spot electricity prices. The analysis of this section provides intuition on the relationship between the amount of reserves from building aggregations and the capacity payments. In practice, estimating the capacity payment is a challenging task that needs to consider additional costs, e.g., due to device wear, but also the competition, i.e., the expected bid prices of generators and/or other load aggregations.

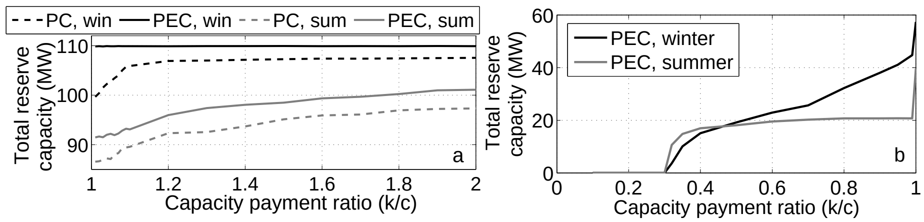

Figure 7: Bid curves of building aggregation in winter and summer. Left: for PC and PEC, and ratios $k/c > 1$. Right: for PEC and ratios $k/c \leq 1$.

6.2 Reserve Product Characteristics

In this section, we fix $k/c = 1.1$ and investigate the influence of important reserve product characteristics on reserve capacities. For PC, we consider reserves with daily or hourly duration, and symmetric or asymmetric capacities. For PEC, we investigate reserves with daily or hourly duration, as well as different (T, ε) pairs, based on Table 1. We summarize simulation results with respect to the total reserve capacity for two weeks in winter and summer in Tables 2 and 3, where positive values denote up- and negative values down-reserves.

For PC and symmetric products[4], hourly reserves increase capacities by 2.9% in winter and 3.3% in summer, compared to daily reserves. If asymmetric reserves are allowed, the aggregation provides significantly more down- than up-reserves in summer, whereas in winter no up-reserves are provided at all. Down-reserves (increasing power consumption) are preferable for buildings equipped with MPC for energy efficient control because the capacity can be offered without increased baseline consumption. The energy efficient control tries to stay as close as possible to the minimum power consumption of the heating/cooling device. In order to provide up-reserves (consumption decrease), a building needs to be able to reduce its consumption without violating occupant comfort. Therefore, the building must schedule its operation (baseline) at a power level higher than the energy optimal, which increases energy costs. On the other hand, down-reserves can be provided while operating at the energy efficient trajectory because the consumption can only increase will tracking the SFC signal.

Compared to PC, PEC increase the reserves for all pairs (T, ε) up to 3.8% in winter and 5.6% in summer. For given T and ε, adopting hourly instead of daily reserves increases the capacities up to 2% in winter and 9.6% in summer. Note that the increase is higher for large T. Notice that shorter T are characterized by smaller εT products, and so constraint (8) on the uncertainty becomes tighter. Therefore, one would intuitively expect that decreasing T increases the reserve capacities monotonically. However, one has to consider that longer T couple more optimization periods, which is why a monotonic behavior is not observed in our

[4]Note that the results of Table 2 are not directly comparable with the values reported in Tables 1 and 2 of [26], because different COP values have been used in the simulations of the two papers.

Table 2: Capacity (MW) of Reserve Products (Power Constraints)

Duration	Symmetry	Winter	Summer
daily	symmetric	±105.9	±89.6
daily	asymmetric	$+0/-196.9$	$+45.3/-195.0$
hourly	symmetric	±109.0	±92.6
hourly	asymmetric	$+0/-213.6$	$+44.8/-198.4$

Table 3: Capacity (MW) of Reserve Products (Power & Energy Constr.)

T (h)	ε (-)	εT (h)	day/win	hour/win	day/sum	hour/sum
1	0.382	0.382	±109.4	±110.0	±92.5	±97.6
2	0.300	0.600	±109.9	±110.7	±93.1	±99.3
4	0.317	1.268	±109.0	±110.2	±92.6	±99.5
6	0.290	1.740	±108.9	±110.4	±93.3	±99.6
8	0.237	1.896	±108.9	±111.0	±94.6	±101.5
12	0.203	2.436	±107.7	±111.1	±92.8	±99.6

simulation results. To gain deeper insight into the dependence of reserve capacities on T and ε, we run simulations for two weeks in winter and summer with $T \in \{1, 2, 4, 6, 8, 12\}$ hours and ε varying from 0.05 to 0.5 with 0.05 increments, i.e., 60 combinations for each season. We assume daily, symmetric reserves and simulate only Lv1, i.e., the reserve scheduling problem, whereas the building HVAC control is not considered. The reason is that many of the combinations of T and ε are not achievable by the Chebyshev filter (20), and so Lv2 will likely be infeasible. However, our analysis provides intuition on the effect of T and ε that could be useful for filter design in a practical application.

In Fig. 8(a), we show the total reserve capacities for each of the simulated cases in winter. Similar results are obtained for summer, but are omitted here due to space limitations. The dependence of reserve capacities on T and ε demonstrates a clear pattern: increasing any of the two parameters reduces the reserve capacity. Note that the capacity is more sensitive to ε than to T. To better illustrate this, we present the results of Fig. 8(a) based on the product εT in Fig. 8(b). As expected, increasing εT generally decreases the reserve capacity. The same εT can be obtained by different (T, ε) pairs: for example, both $(8\text{ h}, 0.25)$ and $(4\text{ h}, 0.5)$ obtain $\varepsilon T = 2$, but the first achieves a capacity 96.72 MW, whereas the second achieves 93.61 MW. In this case, the capacity is more sensitive to ε than to T. We compared all cases in winter and summer with the same εT and found out that in 70% of them smaller ε is preferable to smaller T. This means that buildings can cope easier with signals that are significantly energy constrained over long periods than signals that are moderately constrained over shorter periods.

7 Concluding Remarks

In this paper, we proposed a new framework based on robust optimization and MPC for scheduling and provision of secondary frequency control (SFC) reserves by the HVAC systems of commercial building *aggregations*. The framework incorporates tractable methods to account for energy-constrained SFC signals, and relies on decentralized reserve provision to keep real-time communication requirements low and preserve privacy. We also presented how the framework can be used to estimate the SFC reserve potential from commercial buildings.

Our analysis was based on four main assumptions: (a) there is no plant-model mismatch; (b) the predictions of weather and occupancy are perfect; (c) all building states can be measured; and (d) the reaction of heating/cooling devices is fast and does not cause any wear. Therefore, the reported results provide an upper bound on the amount of reserves from buildings. For a real implementation, additional care must be taken for (a)-(d): accounting for modeling and weather/occupancy prediction errors, use of state estimators, and modeling of the fast dynamics of heating/cooling devices. If multi-zone building models with hundreds of states are available, model reduction techniques [10] can be used to reduce the optimization problem's size. The reserve scheduling problem considered capacity payments, but neglected revenues from reserve energy utilization. Incorporating the latter in the framework is possible if SFC signal scenarios can be generated based on historical data, and is expected to reduce the necessary capacity payments. Additionally, comfort constraints can be relaxed as chance constraints allowing comfort zone violations with a small probability, which is typical for building climate control, while keeping HVAC input constraints robust. Such combination of robust and chance-constrained optimization is likely to increase reserve capacities. In the future, we plan to include these aspects in the framework, implement and test it on a real building.

Overall, our results show that significant amounts of SFC reserves can be *reliably* offered by an aggregation of ~ 100 commercial buildings without loss of occupant comfort. We found that with traditional unconstrained SFC signals asymmetric reserves are preferable for buildings, that energy-constrained SFC signals reduce the necessary capacity payments and increase reserves by up to $\sim 10\%$ compared to traditional SFC signals, and that reducing the duration of reserve products from 1 day to 1 hour increases reserves by up to $\sim 9\%$.

References

[1] Y. Makarov, C. Loutan, J. Ma, P. De Mello, Operational impacts of wind generation on California power systems, IEEE Transactions on Power Systems 24 (2) (2009) 1039–1050.

[2] D. Callaway, I. Hiskens, Achieving controllability of electric loads, Proceedings of the IEEE 99 (1) (2011) 184–199.

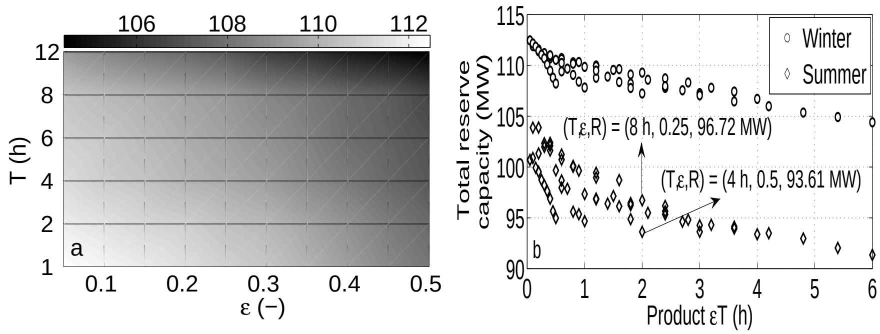

Figure 8: Left: sensitivity of reserve capacity (MW) on averaging period T and bias ε in winter. Right: capacity's dependence on εT in winter and summer.

[3] K. Kalsi, M. Elizondo, J. Fuller, S. Lu, D. Chassin, Development and validation of aggregated models for thermostatic controlled loads with demand response, in: Proceedings of the Hawaii International Conference on Systems Science, Wailea, HI, 2012.

[4] E. Vrettos, S. Koch, G. Andersson, Load frequency control by aggregations of thermally stratified electric water heaters, in: IEEE PES Innovative Smart Grid Technologies Europe (ISGT Europe), Berlin, Germany, 2012.

[5] C. Perfumo, E. Kofman, J. H. Braslavsky, J. K. Ward, Load management: Model-based control of aggregate power for populations of thermostatically controlled loads, Energy Conversion and Management 55 (2012) 36–48.

[6] J. Mathieu, S. Koch, D. Callaway, State estimation and control of electric loads to manage real-time energy imbalance, IEEE Transactions on Power Systems 28 (1) (2013) 430–440.

[7] E. Vrettos, J. L. Mathieu, G. Andersson, Control of thermostatic loads using moving horizon estimation of individual load states, in: Power Systems Computation Conference (PSCC), Wroclaw, Poland, 2014.

[8] F. Oldewurtel, T. Borsche, M. Bucher, P. Fortenbacher, T. Haring, M. G. Vayá, J. L. Mathieu, O. Mégel, E. Vrettos, G. Andersson, A framework for and assessment of demand response and energy storage in power systems, in: IREP IX Bulk Power System Dynamics and Control Symposium, Rethymnon, Greece, 2013.

[9] M. Piette, J. Granderson, M. Wetter, S. Kiliccote, Responsive and intelligent building information and control for low-energy and optimized grid integration, Tech. Rep. LBNL-5662E, Lawrence Berkeley National Laboratory (2012).

[10] D. Sturzenegger, D. Gyalistras, M. Gwerder, C. Sagerschnig, M. Morari, R. Smith, Model predictive control of a Swiss office building, in: REHVA World Congress Clima, Prague, Czech Republic, 2013.

[11] B. Lehmann, D. Gyalistras, M. Gwerder, K. Wirth, S. Carl, Intermediate complexity model for Model Predictive Control of Integrated Room Automation, Energy and Buildings 58 (2013) 250–262.

[12] G. Goddard, J. Klose, S. Backhaus, Model development and identification for fast demand response in commercial HVAC systems, IEEE Transactions on Smart Grid 5 (4) (2014) 2084–2092.

[13] J. E. Braun, Reducing energy costs and peak electrical demand through optimal control of building thermal storage, ASHRAE Transactions 96 (2) (1990) 876–887.

[14] F. Oldewurtel, A. Ulbig, A. Parisio, G. Andersson, M. Morari, Reducing peak electricity demand in building climate control using real-time pricing and model predictive control, in: IEEE Conference on Decision and Control (CDC), Atlanta, USA, 2010.

[15] Y. Ma, F. Borrelli, B. Hencey, B. Coffey, S. Bengea, P. Haves, Model predictive control for the operation of building cooling systems, IEEE Transactions on Control Systems Technology 20 (3) (2012) 797–803.

[16] E. Vrettos, K. Lai, F. Oldewurtel, G. Andersson, Predictive control of buildings for demand response with dynamic day-ahead and real-time prices, in: European Control Conference Conference (ECC), Zurich, Switzerland, 2013.

[17] S. Kiliccote, P. Price, M. Piette, G. Bell, S. Pierson, E. Koch, J. Carnam, H. Pedro, J. Hernandez, A. Chiu, Field testing of automated demand response for integration of renewable resources in California's ancillary services market for regulation products, Tech. Rep. LBNL-5556E, Lawrence Berkeley National Laboratory (2007).

[18] H. Hao, T. Middelkoop, P. Barooah, S. Meyn, How demand response from commercial buildings will provide the regulation needs of the grid, in: 50th Annual Allerton Conference on Communication, Control, and Computing, Monticello, Illinois, USA, 2012.

[19] H. Hao, A. Kowli, Y. Lin, P. Barooah, S. Meyn, Ancillary service for the grid via control of commercial building HVAC systems, in: American Control Conference (ACC), Washington DC, USA, 2013.

[20] Y. Lin, P. Barooah, T. Middelkoop, S. Meyn, Experimental evaluation of frequency regulation from commercial building HVAC system, IEEE Transactions on Smart Grid (to appear).

[21] Y. Lin, P. Barooah, S. Meyn, Low-frequency power grid ancillary services from commercial building HVAC systems, in: IEEE SmartGridComm Conference, Vancouver, Canada, 2013.

[22] Y. Kim, J. L. Kirtley Jr., L. K. Norford, Variable speed heat pump design for frequency regulation through direct load control, in: IEEE T&D Conference and Exposition, Chicago, USA, 2014.

[23] M. Maasoumy, B. M. Sanandaji, A. Sangiovanni-Vincentelli, K. Poolla, Model predictive control of regulation services from commercial buildings to the smart grid, in: American Control Conference (ACC), Portland, USA, 2014.

[24] PJM manual 11: energy & ancillary services market operations, Available online at www.pjm.com/∼/media/documents/manuals/m11.ashx, last accessed 23.01.2015.

[25] P. Barooah, A. Busic, S. Meyn, Spectral decomposition of demand-side flexibility for reliable ancillary services in a smart grid, in: 48th Annual Hawaii International Conference on System Sciences, Kauai, USA, 2015.

[26] E. Vrettos, F. Oldewurtel, F. Zhu, G. Andersson, Robust provision of frequency reserves by office building aggregations, in: World Congress of the International Federation of Automatic Control (IFAC), Cape Town, South Africa, 2014.

[27] Survey on ancillary services procurement and electricity balancing market design, Available online at www.entsoe.eu/publications/position-papers, last accessed 23.01.2015.

[28] Basic principles of ancillary service products, Available online at www.swissgrid.ch, last accessed 23.01.2015.

[29] F. Oldewurtel, Stochastic model predictive control for energy efficient building climate control, Ph.D. thesis, ETH Zürich (2011).

[30] Network code on load-frequency control and reserves, ENTSO-E, Available online at https://www.entsoe.eu/major-projects/network-code-development/load-frequency-control-reserves/Pages/default.aspx, last accessed 07.06.2015.

[31] X. Zhang, G. Schildbach, D. Sturzenegger, M. Morari, Scenario-based MPC for energy-efficient building climate control under weather and occupancy uncertainty, in: IEEE European Control Conference, Zurich, Switzerland, 2013.

[32] E. Camacho, C. Bordon, Model predictive control, Springer, 2013.

[33] J. Löfberg, Automatic robust convex programming, Optimization methods and software 27 (1) (2012) 115–129.

[34] X. Zhang, M. Kamgarpour, P. Goulart, J. Lygeros, Selling robustness margins: A framework for optimizing reserve capacities for linear systems, in: IEEE Conference on Decision and Control, Los Angeles, USA, 2014.

[35] S. P. Boyd, L. Vandenberghe, Convex optimization, Cambridge university press, 2004.

[36] H. Hagiwara, S. Kawachi, J. Baba, K. Furukawa, E. Shimoda, S. Numata, Fundamental study on controllability of heat pump power consumption for fluctuated power compensation, in: IEEE International Conference on Clean Electrical Power, Ischia, Italy, 2011.

[37] I. Avramiotis-Falireas, T. Haring, G. Andersson, M. Zima, Redesign of the automatic generation control scheme in the Swiss power system, in: IEEE PES General Meeting, Washington DC, USA, 2014.

[38] CBE thermal comfort tool, Available online at http://comfort.cbe.berkeley.edu/, last accessed 08.06.2015.

[39] Secondary control power (weekly), Available online at www.swissgrid.ch, last accessed 23.01.2015.